



Office de la propriété
intellectuelle
du Canada

Un organisme
d'Industrie Canada

Canadian
Intellectual Property
Office

An Agency of
Industry Canada



*Bureau canadien
des brevets
Certification*

*Canadian Patent
Office*
RECEIVED
JAN 24 2001
Certification Technology Center 2600

La présente atteste que les documents
ci-joints, dont la liste figure ci-dessous,
sont des copies authentiques des docu-
ments déposés au Bureau des brevets.

This is to certify that the documents
attached hereto and identified below are
true copies of the documents on file in
the Patent Office.

Specification as originally filed, with Application for Patent Serial No: **2,304,542**, on April 10, 2000 by **LORAN NETWORK MANAGEMENT LTD.**, assignee of Nicholas W. Dawes and David Schenkel for "Method of Determining the Route of Packets Trough a Network of Communicating Objects".

S. Paulhus
Agent certificateur/Certifying Officer

October 3, 2000

(Date)

Canada

(CIPO 68)

OPIC



CIPO



2661 0280 #2

IN THE UNITED STATES PATENT AND TRADE MARKS OFFICE

RECEIVED
JAN 24 2001
Technology Center 2600

In re Application of:)
)
Nicholas W. Dawes et al)
)
Serial No: 09/678,783)
)
Filed On: October 4, 2000)
)
For: METHOD OF DETERMINING THE)
ROUTE OF PACKETS THROUGH)
A NETWORK OF COMMUNICATING)
OBJECTS)

Claim for Priority
Under 35 USC 119

November 4, 2000
Our file: 551P08US-1

The Commissioner of Patents
and Trade Marks,
Washington, D.C., 20231,
U.S.A.

Dear Sir:

If any charges or fees must be paid in connection with the
following communication, they may be paid out of deposit account
No. 16-0600.

Enclosed is a copy of the corresponding Canadian application
No. 2,304,542, certified by the Canadian Patent Office and
submitted herewith pursuant to 35 USC 119 and by the
International Convention for Protection of Industrial Properties
and by similar treaties.

Respectfully submitted

PASCAL & ASSOCIATES

by: 
Harold C. Baker - Reg. 19333

Pascal & Associates,
P.O. Box 11121, Station H,
Ottawa, Canada,
K2H 7T8

ABSTRACT

The invention has several parts: the decomposition of the problem of path determination in WAN portions and LAN portions, a novel method for determining the path in WANs, novel methods for determining the path in LANs and finally a novel method for aggregating and reporting the end to end performance of a path, when considered as a whole. A general principle is also applied to computation and presentation of path performance: the performance of a sequence of network elements is a function of the performance of each of the network elements in the sequence. The application of this principle leads to various novel inventions.

Title: Method of determining the route of packets through a network of communicating objects.

Field of the invention:

This invention relates to a method of determining the route of packets through a TCP/IP network of communicating objects.

Description of the invention:

1: Operators of communications networks often need to know the path that packets take between two points in the network. For example, a user may complain that they cannot get through to a server. The operator then wants to determine the path from the user to the server to see if any problem on an element on the path is preventing the user from communicating with the server. With an exact knowledge of the physical topology of the network it may happen that the path from the user to the server is unambiguous (ie: there is only one permitted route). The physical topology can be determined for example by the methods described in US Patents 5,926,462 and 5,933,416 (Dawes et al, ' Method of determining the topology of a network of objects'). However, in many networks the network has redundant or multiple links and therefore the path between some pairs of objects cannot be determined uniquely from the knowledge of the physical topology.

2: Redundant or multiple links can be created between routers (ie: in Wide Area Networks). They can also be created between switches and other devices (ie: in Local Area Networks), especially in switches which support VLANs (Virtual Local Area

Networks). VLANs are a special case of LANs which are especially prone to have ambiguity in paths, if working only from the physical topology. If the entire route between any two objects in the network is decomposed into a series of subpaths, where a subpath is either a WAN or a LAN, the path can be unambiguously determined by applying methods appropriate to WANs or to LANs. These methods are described in the body of this application, but can be summarised as follows.

3: First a method determines for each path to be assessed if that path could be ambiguous. This method can be described in three rules as follows:

- 3.1: a path cannot pass through a broken device or interface.
- 3.2: a path cannot pass through an interface which although up, is carrying no traffic.
- 3.3: a path cannot loop back on itself.

Only paths which are ambiguous need further analysis. This analysis is different from WANs than for LANs.

4: The LAN and WAN segments are decomposed in three steps (which are novel):

- 4.1: by identifying network objects in a topology as being routers or not routers.
- 4.2: by knowing the logical or physical topology of the network, partitioning non-router objects into physically or logically discrete LAN segments which are isolated from other LAN segments by one or more routers.
- 4.3: by knowing the logical or physical topology of the network, partitioning router objects into WAN segments which are isolated from other WAN segments by one or more LAN segments.

5: The determination of a path between two objects through a WAN can be performed by the well known traceroute function, which is widely available. Traceroute returns the sequence of routers in the current path from the source to the destination. However, this allows only the computer performing the traceroute to determine the current path from itself (the source) through routers to other objects (destinations). Unless the traceroute function can be invoked from different sources in the network, a single controlling computer cannot determine the WAN paths between arbitrary points. These other sources have to be computers supporting the traceroute function and must be able to report their results back to a controlling computer or computers, where the aggregation of paths can be performed. One novelty in the method is that these other sources (termed 'beacons') can be optimally placed in the network, if the network's physical topology is known. This drastically reduces the number of beacons and the overhead of the traceroute calls on the network, as well as ensuring complete coverage. Another novelty in the method is that these results are aggregated so that

WAN paths can be added together to describe a longer path. Another novelty in the method is that the WAN paths can be added to LAN paths to produce the path for any arbitrary sequence of WAN and LAN network sections.

6: The determination of the path through a LAN with ambiguous path options requires different analysis than for WANs, because the traceroute function only works for routers (WANs) and a LAN by definition contains no routers. Consider a path from source to destination. The packets from the source flow to the destination. As these packets flow into a port on a switch or hub, that switch or hub will record the MAC address of the source. By reading the tables in switches and hubs that record these MAC addresses, one can determine which ports on which devices 'see' the MAC address of the source. Knowing the topology of the network one can now determine which ports on each device the packets enter and leave on the path from source to destination. Should the source only transit rarely to the destination or indeed never at all, any MAC address from a source in the same subnet as the source will follow the same path (except for the originating port of the source). Finally, should no port on a device on the path ever have reported seeing a MAC from a source in the subnet and yet the path is ambiguous at that point, the port should also be in the same logical subnet in order to pass the information from that source. Note that when a path is broken down into source/LAN/WAN/LAN/destination elements the source MAC address in any LAN after the first is now the MAC of the exit port from exit router in the previous WAN into this LAN.

Determination of the LAN path therefore requires repeated reading of the bridge and source address capture information from devices in the network. This can be accomplished by using SNMP.

Note that these methods for LAN and WAN permit the path from a source to a destination to be different from the path from the destination back to the source.

7: Once the path has been determined between two points in a network, the end to end performance of the path can be determined. This is an application of the general principle that the performance of a sequence of network elements is a function of the performance of each of the network elements in the sequence. For example, by knowing the fraction of packets lost in each element in the path, the overall path loss rate from end to end can be computed. A series of algorithms for overall performance are defined below. These require that the element states be determined (eg: break state, the delay over a line or object and the drop rate through an object). Canadian patent application 2,196,133 (Dawes, 'Method for determining the drop rate, the transit delay and the break state of communications objects) describes how these values may be determined.

Many end to end qualities are the sum, product or logical AND of qualities in the elements on the path. The following four are novel examples of these end to end

qualities.

7.1: A path is available only if all elements on the path are simultaneously available.

7.2: The end to end delay on a path is the sum of delays on each element in the path.

7.3: The jitter (moment to moment change in delay) on a path is sum of jitters in each element in the path.

Define:

7.4: $D(i)$: drop rate on a device i . (eg: $D(1) = 0.12$ means 12% of packets are lost in transit of device 1)

T : the end to end transmission fraction over a path from object 1 - N .

then

$$T = \prod_{i=1..N} (1-D(i))$$

8: Once the end to end qualities of a path have been determined, these can be displayed and reported upon. Very commonly the operators of networks need to know the network performance from any in a set of users to a given server, both now and historically. The following novel method lets the data about the path be determined from time to time, recorded and retrieved for display or further analysis.

Let a set of end points for path analysis be defined, either by the operator directly into a file, files or database, or by importing such data from other sources to a file, files or database or by automatically determining the set of users for each server and determining which of these user's are optimally prototypical (for path analysis) for a set of users and selecting those prototypical users as endpoints and servers as other end points. Let the current path between each pair of end points be determined periodically (eg: every 5 minutes) and the end-end qualities of each path be recorded. Examples of these end to end qualities are the break state (ie: the availability of a path) or the path delay. This data for each path represents the source's communication experience with the destination.

Since a single source can act as prototypical for a set of colocated sources communicating to a single destination, only one path from those colocated sources needs to be determined and measured in order to describe the experience of any of the colocated users.

9: Should a path still remain ambiguous, in perhaps a WAN segment of a path, the

path and its ambiguity can be represented as follows. The ambiguous segment is represented as an icon with perhaps a numeric description of how many alternatives it contains. The unambiguous segments could be represented by appropriate icons. Selecting the ambiguous icon (eg: clicking on it) would provide an expanded view of the ambiguities it represents. This representation can be extended so that multiple ambiguous alternatives in WANs or LANs can be interspersed with unambiguous elements. This representation enormously simplifies the display of such ambiguities. For example, suppose a path had five LAN and WAN segments, and two WAN segments were 5 and 8 fold ambiguous respectively. Full display of all alternative paths would required 40 times as much space.

Should one wish to represent the state of each element by its colour, a surrounding colour or by some other indicator, the worst and best state of each ambiguous segment can also be displayed. A segment's state is for any quality (eg: delay) is just the end-end quality computed from the start object in the section of ambiguous path to the end object in that section. This means that in a representation of the ambiguous segments, the state of the segment can still be presented in the same manner as the unambiguous elements. This is another application of the general principle that the performance of a sequence of network elements is a function of the performance of each of the network elements in the sequence.

10: The optimal location of 'beacons' which perform selected traceroutes can be determined as follows. First note that routers provide the interface between a set of subnets and a WAN: ie: routers interconnect LANs and WANs. Second note that paths need only be determined between predefined sources and destinations.

A beacon is needed in the LAN of a router which satisfies both 8.1 and 8.2:

10.1: Includes one or more sources or destinations in any connected LAN
or

includes one or more paths in any connected LAN (ie: a LAN connecting two WAN segments).

10.2: Has multiple up and active connections out into the WAN.

What is claimed is:

1: a method for determining the path between any two objects in a communications network by:

- (a) decomposing the problem into a sequence of LAN and WAN segments,
- (b) determining the path within each LAN and WAN segment,
- (c) combining the set of paths from each segment to form the complete path.

2: a method as defined in claim 1 where

- (a) the LAN and WAN segments are decomposed by identifying network objects in a topology as being routers or not routers,
- (b) by knowing the logical or physical topology of the network, partitioning non-router objects into physically or logically discrete LAN segments which are isolated from other LAN segments by one or more routers,
- (c) by knowing the logical or physical topology of the network, partitioning router objects into WAN segments which are isolated from other WAN segments by one or more LAN segments.

3: a method as defined in claim 1(b) such that by knowing a priori the physical topology of a network a path between points in a WAN can be determined by requiring three conditions as follows:

- (a): a path cannot pass through a broken device or interface.
- (b): a path cannot pass through an interface which although up, is carrying no traffic.
- (c) a path cannot loop back on itself.

4: a method as defined in claim 1(c) such that by knowing a priori the physical topology of a network a path between points in a LAN can be determined by requiring three conditions as follows:

- (a): a path cannot pass through a broken device or interface.
- (b): a path cannot pass through an interface which although up, is carrying no traffic.
- (c) a path cannot loop back on itself.

5: a method as defined in claim 1(b) where

the WAN paths between required end points are determined by performing traceroutes from beacon objects placed in selected LANs

6: a method as defined in claim 1(c) where

the LAN paths between required end points are determined by reading the bridge or source address or ARP tables from devices and requiring that the path from the source to the destination pass through ports which report the MAC address of the source (or: if in a LAN segment beyond the first LAN, the MAC address of the router port passing this path into this LAN segment)

7: a method as defined in claim 1(c) where

the LAN paths between required end points are determined by reading the logical subnet addresses supported on ports in devices and requiring that the path from the source to the destination pass through ports subnet the required subnet of the destination or next router.

8: a method as defined in claim 1(c) where

the LAN paths between required end points are determined by reading the bridge or source address or ARP tables from devices and requiring that the path from the source or any object in the subnet of the source to the destination or any object in the subnet of the destination pass through ports which report the MAC address of the source (or: if in a LAN segment beyond the first LAN, the MAC address of the router port passing this path into this LAN segment)

9: a method for determining the performance of the path between any two objects in a communications network by first determining the elements in the path and then computing the aggregate performance of the path by a mathematical function of the individual element performance where this mathematical function is

- (a): addition (eg: the sum of delays or jitter)
- (b): multiplication (eg: the product of transit percentage)
- (c): boolean AND (eg: a path is available only if all elements on it are available)

10: a method for determining the optimal location of 'beacons' which perform traceroute or other path determination functions into LANs such that a beacon should be placed in each LAN with one or more routers connected which have multiple up and active connections out into the WAN and that LAN satisfies condition (a) or (b) below:

- (a) :Includes one or more sources or destinations in any connected LAN
- (b): includes one or more paths in any connected LAN (ie: a LAN connecting two WAN segments).

11: a method for representing ambiguous segments in a path by a single object or icon which when selected can indicate details of the alternative paths that it represents.

12: a method for representing ambiguous segments and unambiguous network elements or segments in a path where the segment or element state is indicated by coloured background or surround or other visual signal.

13: a method for representing a path and its alternatives on the map by highlighting elements and segments and letting this highlighting be visually separable or orthogonal from indications of element or network state.

14: a method for determining which pairs of sources and destinations need to have their paths periodically determined, where these pairs are entered either by human entry or by determining prototypical users from their similarity in location and destination to other users (eg: by a clustering technique based on user subnets and destination addresses).

15: a method for determining by claim(1) and by any other claim in this application or by any other method the path between two objects where these objects are one of a set defined as in claim(14) and recording the end to end state of these paths in terms of availability, delay, jitter, packet drop rate and other qualities and the event that is causing maximum impact on these states at that time by knowledge of the network elements on the path and their elemental states.

16: a method of reporting and/or displaying the state of paths at present or historically for paths for predetermined end points (by the method in claim 14 or by any other method)

(a): comparing and reporting the current state against the average, worst and best for the same or related time periods in the past.

(b): comparing and reporting the historical states over a past period against the average, worst and best for the similar previous time periods.

(c): indicating the frequency with which various alternative paths have been in use, when, what the end to end states were active for this different alternatives and for each path and period indicating the network element(s) which contributed to the degradation in state, optionally ranked so that those with greatest impact are indicated first.

17: a method of reporting and/or displaying the state of paths at present or historically for paths for end points defined directly and not known previously by

(a): determining the current path (eg: by claim (1) or any claim in this patent or by any

other method), its current and historical end to end performance and the impact on the current performance of network elements or sets of elements:

(b): comparing and reporting the current state against the average, worst and best for the same or related time periods in the past.

(c): comparing the historical states over a past period against the average, worst and best for the similar previous time periods.

(d): indicating for each period in the past the network element(s) which contributed to the degradation in state, optionally ranked so that those with greatest impact are indicated first.

